

Exam Code: JN0-480

Exam Name: Data Center, Specialist



Exam A

QUESTION 1

What does EVPN use to identify which remote leaf device advertised the EVPN route?

- A. a route distinguisher value
- B. a community tag
- C. a route target value
- D. a VRF target value

Correct Answer: A

Section:

Explanation:

EVPN uses a route distinguisher (RD) value to identify which remote leaf device advertised the EVPN route. An RD is a 64-bit value that is prepended to the EVPN NLRI to create a unique VPNv4 or VPNv6 prefix. The RD value is usually derived from the IP address of the PE that originates the EVPN route. By comparing the RD values of different EVPN routes, a PE can determine which remote PE advertised the route and which VRF the route belongs to. The other options are incorrect because:

B) a community tag is wrong because a community tag is an optional transitive BGP attribute that can be used to group destinations that share some common properties. A community tag does not identify the source of the EVPN route.

C) a route target value is wrong because a route target (RT) value is an extended BGP community that is used to control the import and export of EVPN routes between VRFs. An RT value does not identify the source of the EVPN route.

D) a VRF target value is wrong because there is no such thing as a VRF target value in EVPN. A VRF is a virtual routing and forwarding instance that isolates the IP traffic of different VPNs on a PE. A VRF does not have a target value associated with it. Reference:

EVPN Fundamentals

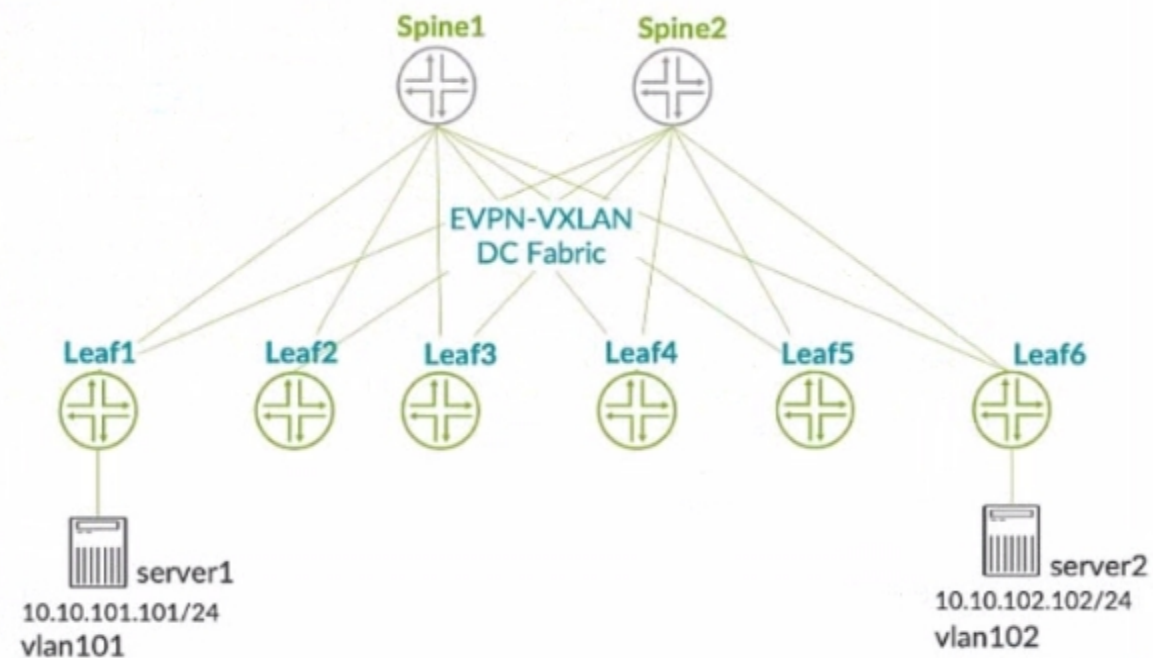
RFC 9136 - IP Prefix Advertisement in Ethernet VPN (EVPN)

EVPN Type-5 Routes: IP Prefix Advertisement

Understanding EVPN Pure Type 5 Routes

QUESTION 2

Exhibit.



You connect two single-homed servers using Juniper Apstra as shown in the exhibit. You are using the ERB design blueprint with two virtual networks in a common routing zone. In this scenario, which two types of VXLAN tunnels will be automatically created by the EVPN control plane? (Choose two.)

- A. EVPN signaled route Type-8 VXLAN tunnels
- B. EVPN signaled route Type-3 VXLAN tunnels
- C. EVPN signaled route Type-6 VXLAN tunnels
- D. EVPN signaled route Type-2 VXLAN tunnels



Correct Answer: B, D

Section:

Explanation:

According to the Juniper documentation¹, EVPN route Type-3 is used to advertise the IP address of the VTEP and the VNIs that it supports. This allows the VTEPs to discover each other and form VXLAN tunnels for the VNIs that they have in common. EVPN route Type-2 is used to advertise the MAC and IP addresses of the hosts connected to the VTEPs. This allows the VTEPs to learn the MAC-to-IP bindings and the MAC-to-VTEP mappings for the hosts in the same VNI. Therefore, these two types of VXLAN tunnels will be automatically created by the EVPN control plane when using Juniper Apstra with the ERB design blueprint and two virtual networks in a common routing zone. Reference: Example: Configure an EVPN-VXLAN Centrally-Routed Bridging Fabric

QUESTION 3

In the Juniper Apstra design phase, which object dictates port count, port speed, and how the ports would be used?

- A. logical devices
- B. rack type
- C. network devices
- D. interface map

Correct Answer: D

Section:

Explanation:

Interface maps are objects that map interfaces between logical devices and physical hardware devices in the Juniper Apstra design phase. They dictate port count, port speed, and how the ports would be used for achieving the intended network configuration rendering. Interface maps also allow you to select device ports, transformations, and interfaces, provision breakout ports, and disable unused ports. For more information, see Interface Maps (Datacenter Design). Reference:

QUESTION 4

You want to keep virtual networks isolated from each other within the Juniper Apstra system. In this scenario, what are three ways to accomplish this task? (Choose three.)

- A. Disable IPv4 connectivity when creating the virtual network within the same Routing Zone.
- B. Enable Security Policy for virtual networks in the same Routing Zone.
- C. Disable Route Target exports when creating the Routing Zones.
- D. Use Connectivity Templates to block access within the same Routing Zone.
- E. Put each network in different Routing Zones.

Correct Answer: B, D, E

Section:

Explanation:

To keep virtual networks isolated from each other within the Juniper Apstra system, you can use one or more of the following methods:

Enable Security Policy for virtual networks in the same Routing Zone. This allows you to define rules that control the traffic flow between different virtual networks within the same routing zone. You can specify the source and destination virtual networks, the protocol, the port, and the action (allow or deny) for each rule. The security policy is applied on the ingress interface of the leaf devices¹.

Use Connectivity Templates to block access within the same Routing Zone. This allows you to customize the connectivity between different racks within the same routing zone. You can create templates that define the link type, the routing protocol, and the access control list (ACL) for each rack pair. The ACL can be used to filter the traffic based on the source and destination IP addresses, the protocol, and the port².

Put each network in different Routing Zones. This allows you to create logical boundaries between different virtual networks based on the route target (RT) values. A routing zone is a collection of virtual networks that share the same RT for importing and exporting routes. Virtual networks in different routing zones do not exchange routes with each other, unless you configure remote EVPN gateways to connect them³. Reference:


Security Policy

Connectivity Templates

Routing Zones

QUESTION 5

Exhibit.



The 10.100.0.0/16 route is being advertised into your BGP IP fabric. ECMP load balancing has been properly enabled on all devices. In this scenario, how many routes will the leaf device in AS 65000 receive for the 10.100.0.0/16 prefix?

- A. 3
- B. 1
- C. 2
- D. 4

Correct Answer: A

Section:

Explanation:

The leaf device in AS 65000 will receive three routes for the 10.100.0.0/16 prefix, one from each spine device in AS 65001, AS 65002, and AS 65003. Since ECMP load balancing is enabled, the leaf device will install all three routes in its routing table and distribute the traffic among them. The other options are incorrect because:

B) 1 is wrong because the leaf device will not receive only one route for the prefix. It will receive multiple routes from different spine devices and use ECMP to load balance among them.

C) 2 is wrong because the leaf device will not receive only two routes for the prefix. It will receive three routes from three spine devices, as explained above.

D) 4 is wrong because the leaf device will not receive four routes for the prefix. It will receive three routes from three spine devices, as explained above. The fourth spine device in AS 65004 is not directly connected to the leaf device and will not advertise the prefix to it. Reference:

IP Fabric Underlay Network Design and Implementation

BGP Multipath load sharing iBGP and eBGP

ECMP Load Balancing

QUESTION 6

Using the Juniper Apstra multitenancy capabilities, which approach will allow a tenant to interconnect two different routing zones?

- A. Interconnection is the default behavior.
- B. Use interconnection through the fabric spine nodes.
- C. Interconnection cannot be enabled.
- D. Use interconnection through an external gateway.



Correct Answer: D

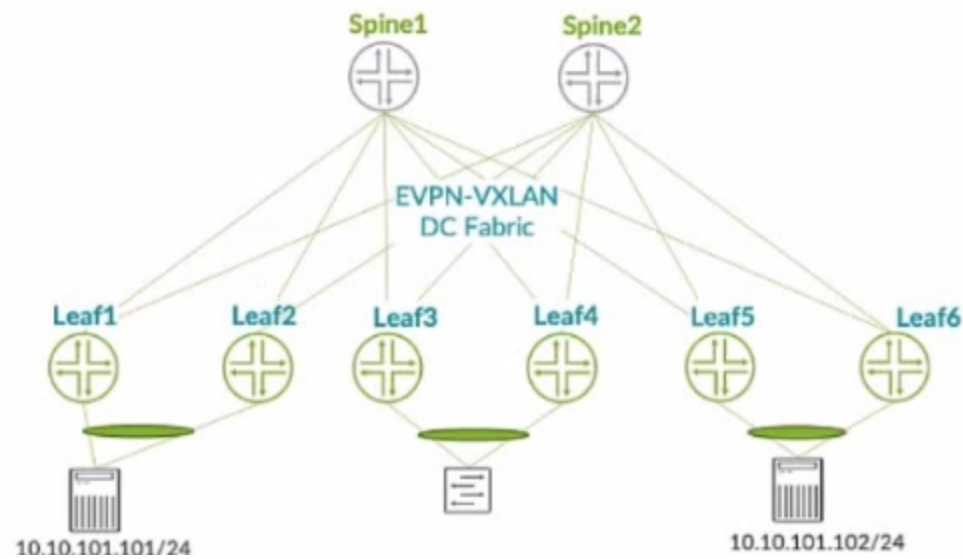
Section:

Explanation:

According to the Juniper documentation¹, a routing zone is an L3 domain, the unit of tenancy in multi-tenant networks. You create routing zones for tenants to isolate their IP traffic from one another, thus enabling tenants to re-use IP subnets. In addition to being in its own VRF, each routing zone can be assigned its own DHCP relay server and external system connections. You can create one or more virtual networks within a routing zone, which means a tenant can stretch its L2 applications across multiple racks within its routing zone. For virtual networks with Layer 3 SVI, the SVI is associated with a Virtual Routing and Forwarding (VRF) instance for each routing zone isolating the virtual network SVI from other virtual network SVIs in other routing zones. If you're using multiple routing zones, external system connections must be from leaf switches in the fabric. Routing between routing zones must be accomplished with external systems. Therefore, the correct answer is D. Use interconnection through an external gateway. Reference: Routing Zones

QUESTION 7

Exhibit.



In the EVPN-VXLAN data center fabric bridged overlay architecture shown in the exhibit, the servers are connected to Leaf1 and Leaf6 using the same virtual network identifier (VNI). Which two statements are correct in this scenario? (Choose two.)

- A. The underlay must use IRB interfaces.
- B. The underlay must be provisioned with PIMv2.
- C. Loopback IPv4 addresses must be advertised into the EBGp underlay from leaf and spine devices.
- D. The underlay EBGp peering's must be established between leaf and spine devices.

Correct Answer: C, D

Section:

Explanation:

In the EVPN-VXLAN data center fabric bridged overlay architecture shown in the exhibit, the servers are connected to Leaf1 and Leaf6 using the same virtual network identifier (VNI). This means that the servers belong to the same Layer 2 domain and can communicate with each other using VXLAN tunnels across the fabric. The underlay network provides the IP connectivity between the leaf and spine devices, and it uses EBGp as the routing protocol. Therefore, the following two statements are correct in this scenario:

Loopback IPv4 addresses must be advertised into the EBGp underlay from leaf and spine devices. This is because the loopback addresses are used as the source and destination IP addresses for the VXLAN tunnels, and they must be reachable by all the devices in the fabric. The loopback addresses are also used as the router IDs and the BGP peer addresses for the EBGp sessions.

The underlay EBGp peering's must be established between leaf and spine devices. This is because the EBGp sessions are used to exchange the underlay routing information and the EVPN routes for the overlay network. The EBGp sessions are established using the loopback addresses of the devices, and they follow a spine-and-leaf topology, where each leaf device peers with all the spine devices, and each spine device peers with all the leaf devices.

The following two statements are incorrect in this scenario:

The underlay must use IRB interfaces. This is not true, because the underlay network does not provide any Layer 3 gateway functionality for the overlay network. The IRB interfaces are used to provide inter-VXLAN routing within the fabric, which is not the case in the bridged overlay architecture. The IRB interfaces are used in the edge-routed bridging (ERB) or the centrally-routed bridging (CRB) architectures, which are different from the bridged overlay architecture.

The underlay must be provisioned with PIMv2. This is not true, because the underlay network does not use multicast for the VXLAN tunnels. The VXLAN tunnels are established using EVPN, which uses BGP to distribute the MAC and IP addresses of the end hosts and the VTEP information of the devices. EVPN eliminates the need for multicast in the underlay network, and it provides optimal forwarding and fast convergence for the overlay network.

Exploring EVPN-VXLAN Overlay Architectures -- Bridged Overlay

EVPN LAGs in EVPN-VXLAN Reference Architectures

EVPN-VXLAN Configuration Guide

QUESTION 8

Exhibit.





You are working to build an ESI-LAG for a multihomed server. The ESI-LAG is not coming up as multihomed. Referring to the exhibit, what are two solutions to this problem? (Choose two.)

- A. The gateway IP addresses on both devices must be different.
- B. The LACP system ID on both devices must be the same.
- C. The loopback IP addresses on both devices must be the same.
- D. The ESI ID on both devices must be the same.

Correct Answer: B, D

Section:

Explanation:

According to the Juniper documentation¹, an ESI-LAG is a link aggregation group (LAG) that spans two or more devices and is identified by an Ethernet segment identifier (ESI). An ESI-LAG provides redundancy and load balancing for a multihomed server in an EVPN-VXLAN network. To configure an ESI-LAG, you need to ensure that the following requirements are met:

The LACP system ID on both devices must be the same. This ensures that the LACP protocol can negotiate the LAG parameters and form a single logical interface for the server.

The ESI ID on both devices must be the same. This ensures that the EVPN control plane can advertise the ESI-LAG as a single Ethernet segment and synchronize the MAC and IP addresses of the server across the devices.

The VLAN ID and VNI on both devices must be the same. This ensures that the server can communicate with other hosts in the same virtual network and that the VXLAN encapsulation and decapsulation can work properly.

In the exhibit, the LACP system ID and the ESI ID on both devices are different, which prevents the ESI-LAG from coming up as multihomed. Therefore, the correct answer is B and D. The LACP system ID on both devices must be the same and the ESI ID on both devices must be the same. Reference: ESI-LAG Made Easier with EZ-LAG, Example: Configuring an ESI on a Logical Interface With EVPN-MPLS Multihoming, Introduction to EVPN LAG Multihoming

QUESTION 9

In the case of IP Clos data center five-stage fabric design, what are two roles of the super spines? (Choose two.)

- A. Super spines are used to interconnect two different data center pods.
- B. Super spines connect to all spine devices within the five-stage architecture.
- C. Super spines are used to connect leaf nodes within a data center pod.
- D. Super spines are always connected to an external data center gateway.

Correct Answer: A, B

Section:

Explanation:

In the case of IP Clos data center five-stage fabric design, the super spines are the devices that provide the highest level of aggregation in the network. They have two main roles:

Super spines are used to interconnect two different data center pods. A pod is a cluster of leaf and spine devices that form a 3-stage Clos topology. A 5-stage Clos topology consists of multiple pods that are connected by the super spines. This allows for scaling the network to support more devices and bandwidth.

Super spines connect to all spine devices within the five-stage architecture. The spine devices are the devices that provide the second level of aggregation in the network. They connect to the leaf devices, which are the devices that provide access to the end hosts. The super spines connect to all the spine devices in the network, regardless of which pod they belong to. This provides any-to-any connectivity between the pods and enables optimal routing and load balancing.

The following two statements are incorrect in this scenario:

Super spines are used to connect leaf nodes within a data center pod. This is not true, because the leaf nodes are connected to the spine nodes within the same pod. The super spines do not connect to the leaf nodes directly, but only through the spine nodes.

Super spines are always connected to an external data center gateway. This is not true, because the super spines are not necessarily involved in the external connectivity of the data center. The external data center gateway is a device that provides the connection to the outside network, such as the Internet or another data center. The external data center gateway can be connected to the super spines, the spine nodes, or the leaf nodes, depending on the design and the requirements of the network.

5-stage Clos Architecture --- Apstra 3.3.0 documentation

5-Stage Clos Architecture | Juniper Networks

Extreme Fabric Automation Administration Guide

QUESTION 10

IBA probes analyze telemetry data from specified devices within a blueprint. Which component Identifies devices that supply data for a specific probe?

- A. data selector
- B. processor
- C. search engine
- D. graph query

Correct Answer: D

Section:

Explanation:

A graph query is a component that identifies devices that supply data for a specific probe. A graph query is an expression that matches nodes in the Apstra graph database based on their attributes, such as device name, role, type, or tag. A graph query can be used to select the source devices for the input processors of a probe, as well as to filter the data by device attributes in the subsequent processors of a probe.¹²Reference:

Probes

Apstra IBA Getting Started Tutorial



QUESTION 11

Which attribute enables Juniper Apstra to scale and manage thousands of devices with a single server instance?

- A. Apstra is installed as a cloud resource.
- B. Apstra is based on NGINX.
- C. Apstra is available as an OVA.
- D. Apstra is a distributed state system.

Correct Answer: D

Section:

Explanation:

The attribute that enables Juniper Apstra to scale and manage thousands of devices with a single server instance is that Apstra is a distributed state system. This means that Apstra uses a graph database to store the network topology and configuration data in a distributed and replicated manner across multiple server nodes. This allows Apstra to handle large-scale networks with high performance, reliability, and availability. Apstra also uses a stateful orchestration engine that ensures the network state is always consistent with the intent of the blueprint, which is the logical representation of the network design and behavior. Apstra can automatically detect and resolve any discrepancies between the desired and actual network state, as well as handle any changes or failures in the network. The other options are incorrect because:

A) Apstra is installed as a cloud resource is wrong because Apstra can be installed either as a cloud resource or as an on-premises resource. Apstra is available as a virtual machine image that can be deployed on various hypervisors, such as VMware ESXi, QEMU/KVM, Microsoft Hyper-V, or Oracle VirtualBox. Apstra can also be deployed on public cloud platforms, such as Amazon Web Services (AWS) or Microsoft Azure. However, the installation method does not affect the scalability of Apstra, which is determined by the distributed state system architecture.

B) Apstra is based on NGINX is wrong because Apstra is not based on NGINX, but on Python and Django. NGINX is a web server and reverse proxy that Apstra uses to serve the web user interface and the REST API. However, NGINX is not the core component of Apstra, and it does not affect the scalability of Apstra, which is determined by the distributed state system architecture.

C) Apstra is available as an OVA is wrong because Apstra is available as an OVF, not an OVA. An OVF (Open Virtualization Format) is a standard format for packaging and distributing virtual machine images. An OVA (Open Virtual Appliance) is a single file that contains the OVF and the virtual disk images. Apstra provides an OVF file that can be imported into various hypervisors, such as VMware ESXi, QEMU/KVM, Microsoft Hyper-V, or Oracle VirtualBox. However, the availability of Apstra as an OVF does not affect the scalability of Apstra, which is determined by the distributed state system architecture. Reference:

QUESTION 12

You have a virtual network that needs controlled access to other virtual networks in the same routing zone. Using the Juniper Apstra UI, which feature would be used to accomplish this task?

- A. interface policy
- B. anti-affinity policy
- C. routing policy
- D. security policy

Correct Answer: D

Section:

Explanation:

A security policy is the feature that would be used to accomplish the task of controlling access to other virtual networks in the same routing zone using the Juniper Apstra UI. A security policy allows you to define rules that specify which traffic is allowed or denied between different virtual networks, IP endpoints, or routing zones. A security policy can be applied to one or more virtual networks in the same routing zone, and it can use various criteria to match the traffic, such as source and destination IP addresses, protocols, ports, or tags. A security policy can also support DHCP relay, which enables the forwarding of DHCP requests from one virtual network to another. The other options are incorrect because:

A) interface policy is wrong because an interface policy is a feature that allows you to configure the interface parameters for the devices in a blueprint, such as interface names, speeds, types, or descriptions. An interface policy does not affect the access control between different virtual networks in the same routing zone.

B) anti-affinity policy is wrong because an anti-affinity policy is a feature that allows you to prevent certain devices or logical devices from being placed in the same rack or leaf pair in a blueprint. An anti-affinity policy is used to enhance the availability and redundancy of the network, not to control the access between different virtual networks in the same routing zone.

C) routing policy is wrong because a routing policy is a feature that allows you to configure the routing parameters for the devices in a blueprint, such as routing protocols, autonomous system numbers, route filters, or route maps. A routing policy does not affect the access control between different virtual networks in the same routing zone, unless the routing policy is used to filter or modify the routes exchanged between different routing zones. Reference:

Security Policy

Interface Policy

Anti-Affinity Policy

Routing Policy

QUESTION 13

What is the purpose of an interface map in Juniper Apstra?

- A. An interface map associates a logical device with a device profile.
- B. An interface map specifies a connection between the interfaces of two devices.
- C. An interface map specifies the number of ports and the port speeds of a logical device
- D. An interface map specifies the connections between racks in a template.

Correct Answer: B

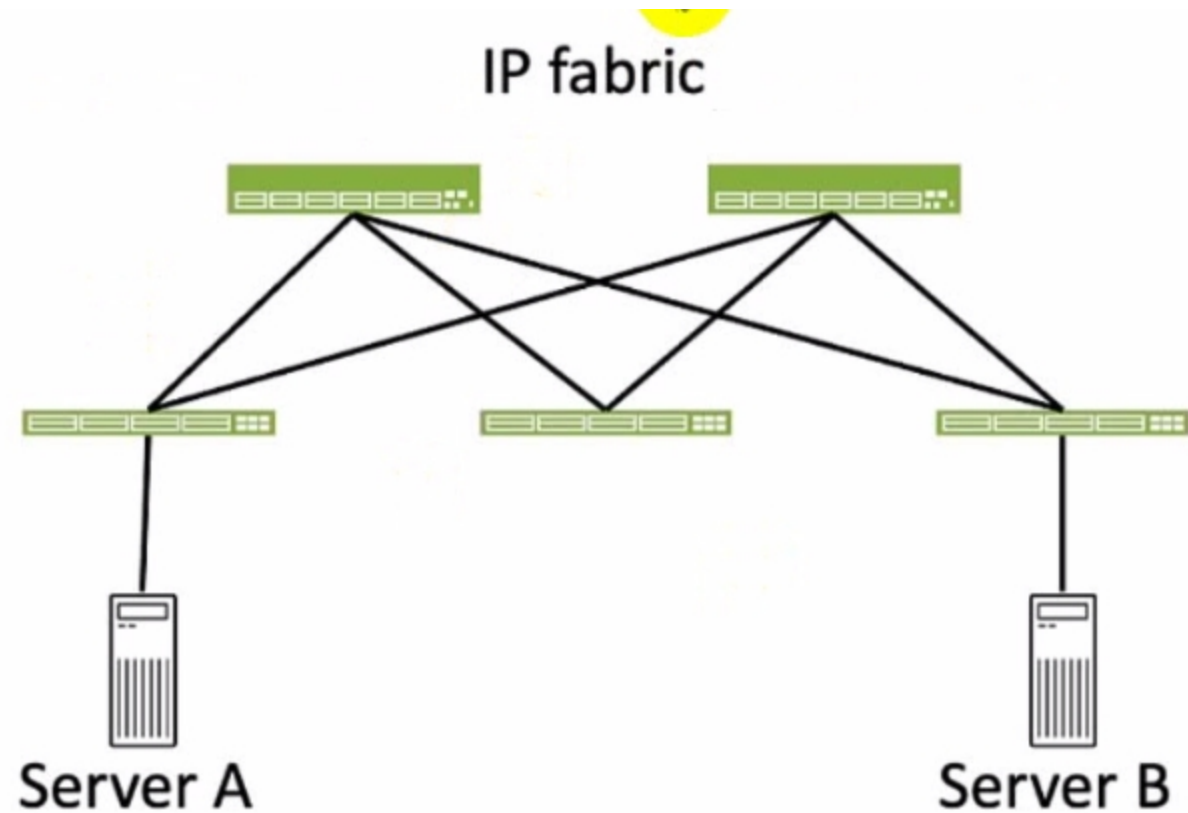
Section:

Explanation:

According to the Juniper documentation¹, an interface map is a configuration template that maps interfaces between logical devices and physical hardware devices (represented with device profiles) while adhering to vendor specifications. An interface map specifies a connection between the interfaces of two devices, such as a leaf and a spine, a leaf and a server, or a leaf and an external gateway. An interface map can also specify port transformations, such as breaking out a 40 GbE port into four 10 GbE ports, or disabling unused ports. An interface map can be used to achieve the intended network configuration rendering and to enable features such as LAG, ESI-LAG, or MLAG. Therefore, the correct answer is B. An interface map specifies a connection between the interfaces of two devices. Reference: Interface Maps (Datacenter Design)

QUESTION 14

Exhibit.



Referring to the exhibit, how many broadcast domains will an Ethernet frame pass through when traversing the IP fabric from Server A to Server B?

- A. 1
- B. 4
- C. 2
- D. 3



Correct Answer: C

Section:

Explanation:

Referring to the exhibit, the image shows a simplified diagram of an IP fabric network connecting two servers, labeled as Server A and Server B. The IP fabric is a network architecture that uses a Clos topology to provide high bandwidth, low latency, and scalability for data center networks. The IP fabric consists of spine and leaf devices that use BGP as the routing protocol and VXLAN as the overlay technology¹.

A broadcast domain is a logical portion of a network where any device can directly transmit broadcast frames to other devices at the data link layer (OSI Layer 2). A broadcast frame is a frame that has a destination MAC address of all ones (FF:FF:FF:FF:FF:FF), which means that it is intended for all devices in the same broadcast domain. A broadcast domain is usually bounded by a router, which does not forward broadcast frames to other networks².

In the exhibit, there are two broadcast domains that an Ethernet frame will pass through when traversing the IP fabric from Server A to Server B. The first broadcast domain is the one that contains Server A and the leaf device that it is connected to. The second broadcast domain is the one that contains Server B and the leaf device that it is connected to. The IP fabric itself is not a broadcast domain, because it uses IP routing and VXLAN encapsulation to transport the Ethernet frames over the Layer 3 network. Therefore, the statement C is correct in this scenario.

The following three statements are incorrect in this scenario:

- A) 1. This is not true, because there are not one, but two broadcast domains that an Ethernet frame will pass through when traversing the IP fabric from Server A to Server B. The IP fabric itself is not a broadcast domain, because it uses IP routing and VXLAN encapsulation to transport the Ethernet frames over the Layer 3 network.
- B) 4. This is not true, because there are not four, but two broadcast domains that an Ethernet frame will pass through when traversing the IP fabric from Server A to Server B. The spine devices and the leaf devices that are not connected to the servers are not part of the broadcast domains, because they use IP routing and VXLAN encapsulation to transport the Ethernet frames over the Layer 3 network.
- D) 3. This is not true, because there are not three, but two broadcast domains that an Ethernet frame will pass through when traversing the IP fabric from Server A to Server B. The IP fabric itself is not a broadcast domain, because it uses IP routing and VXLAN encapsulation to transport the Ethernet frames over the Layer 3 network.

IP Fabric Overview

Broadcast Domain - NetworkLessons.com

QUESTION 15

Which two actions are required during Juniper Apstra's deploy phase? (Choose two.)

- A. Assign device profiles to the blueprint.
- B. Assign user roles to the blueprint.
- C. Assign interlace maps to the blueprint.
- D. Assign resources to the blueprint.

Correct Answer: A, D

Section:

Explanation:

The deploy phase is the final step in the Juniper Apstra data center fabric design and deployment process. In this phase, you apply the Apstra-rendered configuration to the devices and verify the intent of the blueprint. Based on the web search results, we can infer the following actions are required during the deploy phase¹²:

Assign device profiles to the blueprint. This action associates a specific vendor model to each logical device in the blueprint. Device profiles contain extensive hardware model details, such as form factor, ASIC, CPU, RAM, ECMP limit, and supported features. Device profiles also define how configuration is generated, how telemetry commands are rendered, and how configuration is deployed on a device. Device profiles enable the Apstra system to render and deploy the configuration according to the Apstra Reference Design³⁴.

Assign resources to the blueprint. This action allocates the physical devices, IP addresses, VLANs, and ASNs to the logical devices, networks, and routing zones in the blueprint. Resources can be assigned manually or automatically by the Apstra system. Assigning resources ensures that the blueprint has all the necessary elements to generate the configuration and deploy the fabric⁵.

Assign user roles to the blueprint. This action is not required during the deploy phase. User roles are defined at the system level, not at the blueprint level. User roles determine the permissions and access levels of different users in the Apstra system. User roles can be system-defined or custom-defined .

Assign interface maps to the blueprint. This action is not required during the deploy phase. Interface maps are defined at the design phase, not at the deploy phase. Interface maps are objects that map the logical interfaces of a logical device to the physical interfaces of a device profile. Interface maps enable the Apstra system to generate the correct interface configuration for each device in the fabric .Reference:

Deploy

Deploy Device

Device Profiles


Juniper Device Profiles

Resources



QUESTION 16

Exhibit.



Referring to the exhibit, which role does Device A serve in an IP fabric?

- A. leaf
- B. spine
- C. super spine
- D. server

Correct Answer: B

Section:

Explanation:

Device A serves as a spine in an IP fabric. An IP fabric is a network architecture that uses a spine-leaf topology to provide high performance, scalability, and reliability for data center networks. A spine-leaf topology consists of two layers of devices: spine devices and leaf devices. Spine devices are the core devices that interconnect all the leaf devices using equal-cost multipath (ECMP) routing. Leaf devices are the edge devices that connect to the servers, storage, or other network devices. In the exhibit, Device A is connected to four leaf devices using multiple links, which indicates that it is a spine device. The other options are incorrect because:

A) leaf is wrong because a leaf device is an edge device that connects to the servers, storage, or other network devices. In the exhibit, Device A is not connected to any servers, storage, or other network devices, but only to four leaf devices, which indicates that it is not a leaf device.

C) super spine is wrong because a super spine device is a higher-level device that interconnects multiple spine devices in a large-scale IP fabric. A super spine device is typically used when the number of leaf devices exceeds the port density of a single spine device. In the exhibit, Device A is not connected to any other spine devices, but only to four leaf devices, which indicates that it is not a super spine device.

D) server is wrong because a server device is a compute or storage device that connects to a leaf device in an IP fabric. A server device is typically the end host that provides or consumes data in the network. In the exhibit, Device A is not connected to any leaf devices, but only to four leaf devices, which indicates that it is not a server device. Reference:

IP Fabric Underlay Network Design and Implementation

IP Fabric Overview

IP Fabric Architecture

QUESTION 17

A member of your organization made changes to a predefined interface map using Juniper Apstra.

Which two statements are correct in this scenario? (Choose two.)

- A. Changes to interface maps in the global catalog do not affect interface maps that have already been imported into blueprint catalogs
- B. Any changes made to predefined interface maps are discarded when Apstra is upgraded.
- C. Changes made to predefined interface maps will not have an impact on the Apstra software.
- D. Changes to interface maps in the global catalog will raise anomalies that may need to be addressed at the next commit.

Correct Answer: A, B

Section:

Explanation:

According to the Juniper documentation¹, an interface map is a configuration template that maps interfaces between logical devices and physical hardware devices (represented with device profiles) while adhering to vendor specifications. An interface map can be either predefined or custom. A predefined interface map is one that ships with Apstra software and supports most qualified Juniper devices. A custom interface map is one that is created by the user to meet specific requirements. An interface map can be stored in either the global catalog or the blueprint catalog. The global catalog contains all the interface maps that are available for use in any blueprint. The blueprint catalog contains the interface maps that are imported from the global catalog and used in a specific blueprint.

When a member of your organization makes changes to a predefined interface map, the following statements are correct:

Changes to interface maps in the global catalog do not affect interface maps that have already been imported into blueprint catalogs. This means that the existing blueprints that use the original version of the interface map will not be impacted by the changes. However, if you want to use the updated version of the interface map in a new or existing blueprint, you need to import it again from the global catalog.


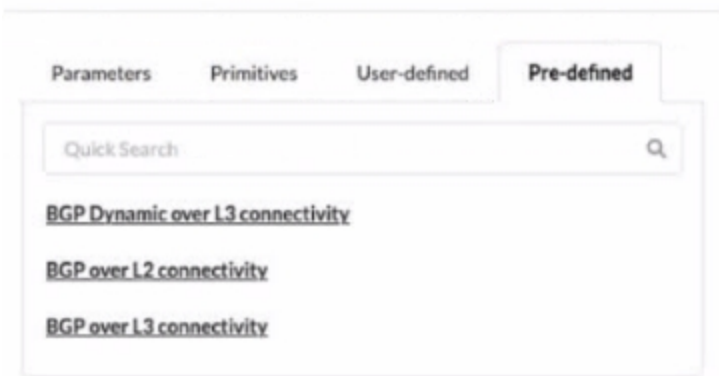
Any changes made to predefined interface maps are discarded when Apstra is upgraded. This means that the changes will not be preserved across different versions of Apstra software. If you want to retain a customized interface map through Apstra upgrades, you need to clone the predefined interface map, give it a unique name, and customize it instead of changing the predefined one directly.

Therefore, the correct answer is A and B. Changes to interface maps in the global catalog do not affect interface maps that have already been imported into blueprint catalogs and any changes made to predefined interface maps are discarded when Apstra is upgraded. Reference: Edit Interface Map | Apstra 4.2 | Juniper Networks

QUESTION 18

Exhibit.

Create Connectivity Template



Referring to the exhibit, which statement is correct?

- A. The gray-solid primitives indicate further configuration is required.
- B. The gray-solid primitives indicate that they are incompatible with the connectivity template design.
- C. The red-striped primitives indicate that they are incompatible with the connectivity template design.
- D. The red-striped primitives indicate that further configuration is required.

Correct Answer: D

Section:

Explanation:

A connectivity template is a set of configuration parameters that can be applied to a device or a group of devices in a blueprint. A blueprint is a logical representation of the network design and intent. A primitive is a basic unit of configuration that can be added to a connectivity template. A primitive can be a link, a peering, a policy, or a service. In the exhibit, the red-striped primitives indicate that further configuration is required for them to be compatible with the connectivity template design. The red stripes mean that the primitive is incomplete or invalid, and it needs to be edited or deleted. For example, the IP Link primitive needs to have the interface name and IP address specified for each end of the link. The other options are incorrect because:

- A) The gray-solid primitives indicate further configuration is required is wrong because the gray-solid primitives indicate that they are compatible with the connectivity template design. The gray color means that the primitive is valid and complete, and it does not need any further configuration.
- B) The gray-solid primitives indicate that they are incompatible with the connectivity template design is wrong because the gray-solid primitives indicate that they are compatible with the connectivity template design, as explained above.
- C) The red-striped primitives indicate that they are incompatible with the connectivity template design is wrong because the red-striped primitives indicate that further configuration is required, not that they are incompatible. The red stripes mean that the primitive is incomplete or invalid, but it can be fixed by editing or deleting it.

Connectivity Templates

Data Center Automation Using Juniper Apstra

Config Rendering in Juniper Apstra

QUESTION 19

What are two system-defined user roles that are available in Juniper Apstra? (Choose two.)

- A. authorized
- B. root
- C. viewer
- D. user

Correct Answer: C, D

Section:

Explanation:

Juniper Apstra provides four system-defined user roles that are available in the Apstra GUI environment. They are: administrator, device_ztp, viewer, and user1. Based on the web search results, we can infer the following statements:

viewer: This role includes permissions to only view various elements in the Apstra system, such as blueprints, devices, design, resources, external systems, platform, and others. Users with this role cannot create, edit, or delete any element¹².

user: This role includes permissions to view and edit various elements in the Apstra system, such as blueprints, devices, design, resources, external systems, platform, and others. Users with this role cannot create or delete any element¹².

authorized: This is not a system-defined user role in Juniper Apstra. It is a term used to describe users who have been authenticated by an external system, such as LDAP, Active Directory, TACACS+, or RADIUS³.

root: This is not a system-defined user role in Juniper Apstra. It is a term used to describe the superuser account on a Linux system, which has full access to all commands and files. Creating a user in the Apstra GUI does not provide that user access to the Apstra platform via SSH. To access the Apstra platform via SSH, you must create a local Linux system user⁴. Reference:

User / Role Management Introduction

User/Role Management (Platform)

AAA Providers

User Profile Management

QUESTION 20

Which two statements are correct about repairing a Juniper Apstra cabling map before deploying your blueprint? (Choose two.)

- A. You must manually change the cabling map to update spine-to-leaf fabric links.
- B. Apstra can use LLDP data from the spine-to-leaf fabric devices to update the connections in the cabling map.
- C. Apstra can use LLDP data from the leaf devices to update the leaf-to-generic connections in the cabling map.
- D. You must manually change the cabling map to update leaf-to-generic links.

Correct Answer: B, C

Section:

Explanation:

The cabling map is a graphical representation of the physical connections between the devices in the data center fabric. It shows the status of the cables, interfaces, and BGP sessions for each device. You can use the cabling map to verify and repair the cabling before deploying your blueprint. Based on the web search results, we can infer the following statements:

Apstra can use LLDP data from the spine-to-leaf fabric devices to update the connections in the cabling map. This is true because Apstra can collect LLDP data from the devices using the Generic Graph Collector processor and use it to update the cabling map automatically. LLDP is a protocol that allows devices to exchange information about their identity, capabilities, and neighbors¹².

Apstra can use LLDP data from the leaf devices to update the leaf-to-generic connections in the cabling map. This is true because Apstra can also collect LLDP data from the leaf devices and use it to update the connections to the generic devices, such as routers, firewalls, or servers. Generic devices are devices that are not managed by Apstra but are part of the data center fabric²³.

You must manually change the cabling map to update spine-to-leaf fabric links. This is false because Apstra can use LLDP data to update the spine-to-leaf fabric links automatically, as explained above. However, you can also manually change the cabling map to override the Apstra-generated cabling, if needed²⁴.

You must manually change the cabling map to update leaf-to-generic links. This is false because Apstra can use LLDP data to update the leaf-to-generic links automatically, as explained above. However, you can also manually change the cabling map to override the Apstra-generated cabling, if needed²⁴. Reference:

LLDP Overview

Edit Cabling Map (Datacenter)

Generic Devices

Import / Export Cabling Map (Datacenter)

QUESTION 21

You are working with a three-stage IP fabric using EBGP for peering.

In this scenario, which two actions are required to implement ECMP? (Choose two.)

- A. Use a load balancing policy applied to the forwarding table as an export policy.
- B. Use a load balancing policy applied to BGP as an export policy.
- C. Use the multipath multiple-as BGP parameter.
- D. Use a load balancing policy applied to BGP as an import policy.



Correct Answer: B, C

Section:

Explanation:

To implement ECMP in IP fabric using EBG, you need to enable BGP to install multiple equal-cost paths in the routing table and to advertise them to the peers. The following actions are required to achieve this:

B) Use a load balancing policy applied to BGP as an export policy. This is true because you need to apply a load balancing policy to BGP as an export policy to allow BGP to advertise multiple paths to the same destination to the peers. By default, BGP only advertises the best path to the peers, which prevents ECMP. A load balancing policy can be configured to match the desired routes and set the multipath attribute to true. This will enable BGP to advertise up to the maximum number of paths configured by the maximum-paths command. For example, the following configuration applies a load balancing policy to BGP as an export policy for the neighbor 10.10.10.1:

```
policy-statement load-balance { term 1 { from { route-filter 192.168.0.0/16 exact; } then { multipath; accept; } } } protocols { bgp { group ebgp { type external; neighbor 10.10.10.1 { export load-balance; } } }
```

C) Use the multipath multiple-as BGP parameter. This is true because you need to enable the multipath multiple-as BGP parameter to allow BGP to install multiple paths from different autonomous systems in the routing table. By default, BGP only installs multiple paths from the same autonomous system, which limits ECMP. The multipath multiple-as parameter can be configured under the BGP group or neighbor level. This will enable BGP to install up to the maximum number of paths configured by the maximum-paths command. For example, the following configuration enables the multipath multiple-as parameter for the BGP group ebgp:

```
protocols { bgp { group ebgp { type external; multipath multiple-as; } }
```

The following options are incorrect because:

A) Use a load balancing policy applied to the forwarding table as an export policy is wrong because applying a load balancing policy to the forwarding table does not affect the BGP advertisement or installation of multiple paths. A load balancing policy applied to the forwarding table only affects how the traffic is distributed among the multiple paths in the forwarding table. It does not enable ECMP in BGP.

D) Use a load balancing policy applied to BGP as an import policy is wrong because applying a load balancing policy to BGP as an import policy does not affect the BGP advertisement of multiple paths. A load balancing policy applied to BGP as an import policy only affects how the BGP routes are accepted or rejected from the peers. It does not enable ECMP in BGP. Reference:

IP Fabric Underlay Network Design and Implementation

Use ECMP to distribute traffic between two paths, one learned by eBGP and one learned by iBGP on a Cisco NX-OS switch

Example: Configure an EVPN-VXLAN Centrally-Routed Bridging Fabric Using EBG

QUESTION 22

What is the purpose of using a routing zone inside Juniper Apstra software?

- A. A routing zone is used to enable L4-L7 inspection inside the fabric.
- B. A routing zone is defined to secure the routing protocols.
- C. A routing zone defined at the Apstra manager level requires firewalls to be deployed.
- D. A routing zone is used to enable the communication between two VNIs within a VRF.



Correct Answer: D

Section:

Explanation:

According to the Juniper documentation¹, a routing zone is an L3 domain, the unit of tenancy in multi-tenant networks. You create routing zones for tenants to isolate their IP traffic from one another, thus enabling tenants to re-use IP subnets. In addition to being in its own VRF, each routing zone can be assigned its own DHCP relay server and external system connections. You can create one or more virtual networks within a routing zone, which means a tenant can stretch its L2 applications across multiple racks within its routing zone. For virtual networks with Layer 3 SVI, the SVI is associated with a Virtual Routing and Forwarding (VRF) instance for each routing zone isolating the virtual network SVI from other virtual network SVIs in other routing zones. Therefore, the correct answer is D. A routing zone is used to enable the communication between two VNIs within a VRF. A routing zone is not used for L4-L7 inspection, securing routing protocols, or requiring firewalls. Those are not the purposes of a routing zone in Juniper Apstra software. Reference: Routing Zones

QUESTION 23

Exhibit.

Filter selected by all selected only unselected only

<input type="checkbox"/>	Name ↕	Rack Type ↕
<input type="checkbox"/>	borderleaf_001	BorderLeaf 2022-03-02 09:33
<input type="checkbox"/>	serverrack_001	ServerRack 2022-03-01 10:19
<input type="checkbox"/>	serverrack_002	ServerRack 2022-03-01 10:19
<input type="checkbox"/>	serverrack_003	ServerRack 2022-03-01 10:19
<input type="checkbox"/>	serverrack_004	ServerRack 2022-03-01 10:19
<input type="checkbox"/>	serverrack_005	ServerRack 2022-03-01 10:19
<input type="checkbox"/>	serverrack_006	ServerRack 2022-03-01 10:19

Referring to the exhibit, how many rack types are used in the staged blueprint?

- A. six
- B. three
- C. seven
- D. two

Correct Answer: D

Section:

Explanation:

Referring to the exhibit, the image shows the Racks table under the Staged menu in the Juniper Apstra UI. The Racks table displays the details of the racks that are used in the blueprint, such as the name, rack type, and date. The rack type is a resource that defines the type and number of leaf devices, access switches, and/or generic systems that are used in rack builds¹. The image shows seven racks in the table, but only two rack types: BorderLeaf and ServerRack. Therefore, the statement D is correct in this scenario.

The following three statements are incorrect in this scenario:

- A) six. This is not true, because there are not six rack types in the table, but only two. The number six corresponds to the number of racks that have the same rack type: ServerRack.
- B) three. This is not true, because there are not three rack types in the table, but only two. The number three does not correspond to any relevant information in the table or the image.
- C) seven. This is not true, because there are not seven rack types in the table, but only two. The number seven corresponds to the total number of racks in the table, not the rack types.

Rack Types (Datacenter Design)

Racks (Staged)

